Examiner-Initiated Interview Summary	Application No. 10/751,016	Applicant(s) RAMESH ET AL.
	Examiner Usmaan Saeed	Art Unit 2166

All Participants:

(1) _Usmaan Saeed (PTO)_.

(2) _Howard L. Speight, Registration No. 37,733_.

Status of Application: _____

(3) _____.

(4) _____.

Date of Interview: _28 January 2008_ **Time:** _____

Type of Interview:

☒ Telephonic
☐ Video Conference
☐ Personal (Copy given to: ☐ Applicant ☐ Applicant's representative)

Exhibit Shown or Demonstrated: ☐ Yes ☒ No
If Yes, provide a brief description: . .

Part I.

Rejection(s) discussed:
None

Claims discussed:
1, 15 and 29

Prior art documents discussed:
None

Part II.

SUBSTANCE OF INTERVIEW DESCRIBING THE GENERAL NATURE OF WHAT WAS DISCUSSED:
See Continuation Sheet

Part III.

☒ It is not necessary for applicant to provide a separate record of the substance of the interview, since the interview directly resulted in the allowance of the application. The examiner will provide a written summary of the substance of the interview in the Notice of Allowability.

☐ It is not necessary for applicant to provide a separate record of the substance of the interview, since the interview did not result in resolution of all issues. A brief summary by the examiner appears in Part II above.

_____
(Examiner/SPE Signature)

_____
(Applicant/Applicant's Representative Signature – if appropriate)

Continuation of Substance of Interview including description of the general nature of what was discussed: A telephone call was made to applicant's representative about the potential amendment in order to allow the case. The representative agreed with the examiner's proposal and gave authorization for examiner's amendment via email containing examiner's proposed amendment. A copy of the amended claims is attached.

Proposed Amendment to Claims

1.   (**currently amended**) A method for representing statistics about a table including one or more rows, each row including a respective value, the method including:

creating one or more histogram buckets, each histogram bucket including a width representing a respective range of values and a height representing a count of rows in the table having values in the range of values;

creating one or more high-bias buckets, each high-bias bucket including one or more high-bias values up to a maximum number of high-bias values (F) that appear in a minimum percentage of rows in the table and for each high-bias value a number of rows that contain the high-bias value;

where the minimum percentage of rows is computed using F and B, where B is the total number of buckets;

repeating the following:

(a) determining an average height of the histogram buckets;

(b) determining a reclassification threshold based on the average height of the histogram buckets; and

(c) concluding that a value associated with one of the one or more histogram buckets occurs in more rows of the table than the reclassification threshold, and, in response, concluding that the number of high-bias values associated with at least one of the one or more high-bias buckets has not reached the maximum number of high-bias values, and, in response, including the value in one of the high-bias buckets for which the number of high-bias values has not reached the maximum number of high-bias values;

until no values included in any of the ranges of values associated with the histogram buckets occur in more than the reclassification threshold number of rows in the table; and

saving in a memory the width and the height of each of the one or more histogram buckets and the one or more high-bias values and numbers of rows for each of the one or more high-bias buckets.

1

2.   (**previously presented**)   The method of claim 1, where the sum of the number of histogram buckets and the number of high-bias buckets is less than a predetermined number.

3.   (**cancelled**)

4.   (**previously presented**)   The method of claim 1, where the reclassification threshold is equal to the average height of the histogram buckets multiplied by (1+S), where S is a positive percentage represented as a decimal.

5-9.   (**cancelled**)

10.   (**currently amended**)   The method of claim 1, further including setting the minimum percentage of rows to $\dfrac{1}{FB}$% , ~~where B is equal to the total number of buckets~~.

11.   (**currently amended**)   The method of claim 1, further including adjusting the minimum percentage of rows to $\dfrac{V(FB-I)}{FB}$% , where ~~B is equal to the total number of buckets,~~ V is equal to the minimum percentage of rows, and I is equal to a number of values represented in high-bias buckets.

12-14.   (**cancelled**)

15.   (**currently amended**) A database system including:
    a massively parallel processing system including:
        one or more nodes;
        a plurality of CPUs, each of the one or more nodes providing access to one or more CPUs;
        a plurality of data storage facilities each of the one or more CPUs providing access to one or more data storage facilities;
        P partitions, each partition residing on one or more data storage facilities;
    a process for representing statistics, where the database system represents statistics about a table including one or more rows, each row including a respective value, the process including:

creating one or more histogram buckets, each histogram bucket including a width representing a respective range of values and a height representing a count of rows in the table having values in the range of values;

creating one or more high-bias buckets, each high-bias bucket including one or more high-bias values up to a maximum number of high-bias values (F) that appear in a minimum percentage of rows in the table and for each high-bias value a number of rows that contain the high-bias value;

where the minimum percentage of rows is computed using F and B, where B is the total number of buckets;

repeating the following:

    (a) determining an average height of the histogram buckets;

    (b) determining a reclassification threshold based on the average height of the histogram buckets; and

    (c) concluding that a value associated with one of the one or more histogram buckets occurs in more rows of the table than the reclassification threshold, and, in response, concluding that the number of high-bias values associated with at least one of the one or more high-bias buckets has not reached the maximum number of high-bias values, and, in response, including the value in one of the high-bias buckets for which the number of high-bias values has not reached the maximum number of high-bias values;

until no values included in any of the ranges of values associated with the histogram buckets occur in more than the reclassification threshold number of rows in the table; and

saving in a memory the width and the height of each of the one or more histogram buckets and the one or more high-bias values and numbers of rows for each of the one or more high-bias buckets.

3

16.   (**previously presented**) The database system of claim 15, where the sum of the number of histogram buckets and the number of high-bias buckets is less than a predetermined number.

17.   (**cancelled**)

18.   (**previously presented**) The database system of claim 15, where the reclassification threshold is equal to the average height of the histogram buckets multiplied by (1+S), where S is a positive percentage represented as a decimal.

19-23. (**cancelled**)

24.   (**currently amended**) The database system of claim 15, further including setting the minimum percentage of rows to $\dfrac{1}{FB}$% ~~,where B is equal to the total number of buckets~~.

25.   (**currently amended**) The database system of claim 15, further including adjusting the minimum percentage of rows to $\dfrac{V(FB-I)}{FB}$%, where ~~B is equal to the total number of buckets,~~ V is equal to the minimum percentage of rows, and I is equal to a number of values represented in high-bias buckets.

26-28.   (**cancelled**)

29.   (**currently amended**) A computer program, stored on a tangible storage medium, for use in representing statistics in a database running in a partitioned parallel environment including P partitions, each partition residing on one or more parallel processing systems, the database including a first table including one or more rows stored in one or more of the P partitions, the program including executable instructions that cause a computer to:

   represent statistics about a table including one or more rows, each row including one or more values, the program further causing the computer to:

   create one or more histogram buckets, each histogram bucket including a width representing a respective range of values and a height representing a count of rows in the table having values in the range of values;

create one or more high-bias buckets, each high-bias bucket including one or more high-bias values up to a maximum number of high-bias values (F) that appear in a minimum percentage of rows in the table and for each high-bias value a number of rows that contain the high-bias value;

where the minimum percentage of rows is computed using F and B, where B is the total number of buckets;

repeat the following:

    (a) determine an average height of the histogram buckets;

    (b) determine a reclassification threshold based on the average height of the histogram buckets; and

    (c) conclude that a value associated with one of the one or more histogram buckets occurs in more rows of the table than the reclassification threshold, and, in response, conclude that the number of high-bias values associated with at least one of the one or more high-bias buckets has not reached the maximum number of high-bias values, and, in response, include the value in one of the high-bias buckets for which the number of high-bias values has not reached the maximum number of high-bias values;

until no values included in any of the ranges of values associated with the histogram buckets occur in more than the reclassification threshold number of rows in the table; and

save in a memory the width and the height of each of the one or more histogram buckets and the one or more high-bias values and numbers of rows for each of the one or more high-bias buckets.

5

30. **(previously presented)** The computer program of claim 29, where the sum of the number of histogram buckets and the number of high-bias buckets is less than a predetermined number.

31. **(cancelled)**

32. **(previously presented)** The computer program of claim 29, where the reclassification threshold is equal to the average height of the histogram buckets multiplied by (1+S), where S is a positive percentage represented as a decimal.

33-37. **(cancelled)**

38. **(currently amended)** The computer program of claim 29, further including executable instructions that cause the computer to set the minimum percentage of rows to $\dfrac{1}{FB}\%$ ~~, where B is equal to the total number of buckets~~.

39. **(currently amended)** The computer program of claim 29, further including executable instructions that cause the computer to adjust the minimum percentage of rows to $\dfrac{V(FB-I)}{FB}\%$, where ~~B is equal to the total number of buckets,~~ V is equal to the minimum percentage of rows, and I is equal to a number of values represented in high-bias buckets.

40-42. **(cancelled)**